

Amendments to the Claims:

This listing of claims replaces all prior versions and listings of claims in the application:

Listing of Claims:

1. (Currently Amended) A method by which more than one client program connected to a network stores the same data item at the same location or locations in a data repository connected to the network, the method comprising:

having a first client program deposit a data item in a data repository, the depositing including

determining a digital fingerprint from the data item using a reproducible pseudorandom process that produces digital fingerprints having a pseudorandom distribution;

storing the data item in the data repository at a physical location or locations associated with the digital fingerprint,

having a second client program initiate a process for depositing a second data item in the data repository, the process including

determining a digital fingerprint from the second data item using the reproducible pseudorandom process;

comparing the digital fingerprint from the second data item to digital fingerprints for data items already stored in the data repository, and determining from the comparing of digital fingerprints, without comparing the entire contents of the second data item to the entire contents of a data item already stored, whether a data item identical to the second data item is already stored in the data repository; and

storing the second data item in the data repository if comparing establishes that a data item identical to the second data item is not already stored in the data repository, and not storing the second data item in the data repository if comparing establishes that a data item identical to the second data item is already stored in the data repository;

wherein the reproducible pseudorandom process produces a digital fingerprint designed to probabilistically guarantee to provide a unique digital fingerprint for every distinct data item sent to the data repository;

wherein ~~the pseudorandom distribution of the digital fingerprints introduces pseudorandomness into~~ the physical locations at which data items are stored in the data repository are determined at least in part by the digital fingerprints.

2-3. (Canceled).

4. (Previously Presented) The method of claim 154 wherein the encrypting of the data item is performed by the client prior to transmitting the data item to the data repository.

5. (Previously Presented) The method of claim 4 further comprising encrypting the key and storing the encrypted key in the data repository.

6. (Original) The method of claim 5 wherein a client or user specific key is used to encrypt the key derived from the content of the data item.

7. (Previously Presented) The method of claim 154 wherein the key derived from the content of the data item is the same for all instances of the data item stored in the repository.

8. (Previously Presented) The method of claim 154 wherein users of the method are grouped into families, and the key derived from the content of the data item is the same for all instances of the data item stored in the repository by users in the same family, but may be different for users in different families.

9. (Previously Presented) The method of claim 1 wherein one or more additional copies or other forms of redundant information about the data items is stored in the data repository for data integrity, availability, or accessibility purposes and not to provide separate storage of the data item for different client programs.

10. (Original) The method of claim 1 further comprising associating the data item with each of a plurality of access-authorization credentials, each of which is uniquely associated with a particular user or client program.

11. (Canceled).

12. (Original) The method of claim 10 wherein the associating of the data item with each of a plurality of access-authorization credentials comprises storing a plurality of named objects, each named object comprising information representative of the data item paired with information representative of one of the access-authorization credentials.

13. (Original) The method of claim 12 wherein the information representative of the data item is a digital fingerprint.

14. (Original) The method of claim 12 wherein the information representative of the access-authorization credential is a cryptographic hash of all or part of the access-authorization credential.

15. (Original) The method of claim 14 wherein the cryptographic hash is an access identifier that uniquely identifies the data item for a particular user or client program.

16. (Original) The method of claim 12 wherein the named object is a data structure created by the client program.

17. (Original) The method of claim 12 wherein the named object is a data structure created by a server program acting on behalf of the repository.

18. (Previously Presented) The method of claim 12 further comprising a client replacing an existing version of a named object with a new version of that named object, by replacing the existing association with a data item stored in the data repository with a new association.

19. (Previously Presented) The method of claim 12 further comprising a client retrieving a data item by accessing a named object using an access-authorization credential to select the named object, and using the contents of the named object to determine the location of the data item in the data repository.

20. (Original) The method of claim 12 wherein the named objects further comprise version information associating different data items with different versions of the named object.

21. (Previously Presented) The method of claim 20 wherein a backup of data items stored in the data repository is accomplished by preserving copies of the current versions of named objects in existence at the time of the backup.

22. (Original) The method of claim 1 wherein records are kept of the association between data items and names in order to define named objects, and wherein data items recorded as being associated with named objects are not deleted from the repository, and wherein named objects are backed up by preserving copies of the named object records in existence at the time of the backup.

23. (Original) The method of claim 21 or 22 wherein a plurality of backups are made at spaced time intervals.

24. (Original) The method of claim 21 or 22 wherein the backup is accomplished by declaring that after a prescribed moment in time a new version of each named object will be created the first time that a new data item is associated with it.

25. (Original) The method of claim 24 wherein the prescribed moment in time is determined separately for each named object.

26. (Original) The method of claim 22 wherein named objects are preserved by creating a new version of each named object each time that a new data item is associated with it.

27. (Original) The method of claim 26, wherein versions of named objects that are deemed unnecessary are deleted.

28. (Original) The method of claim 27, wherein the determination of which versions of a named object to delete is based in whole or in part on the times at which the versions were created, and the intervals between these times.

29. (Original) The method of claim 20 further comprising preparing a digital time stamp of a plurality of named objects to allow a property of these named objects to be proven at a later date.

30. (Original) The method of claim 29 wherein a random or other difficult to guess element is incorporated into the time stamp hash for each named object, to prevent the property from being proven if this element is deleted.

31. (Previously Presented) The method of claim 12 further comprising determining that a data item stored in the data repository is not referenced by any named object, and reusing the storage space used to store the unreferenced data item.

32. (Original) The method of claim 12 further comprising altering one or more properties or parameters associated with an access-authorization credential to change the access rights of a client or user to the data item referenced by that credential.

33. (Currently Amended) The method of claim 1 further comprising a challenge step to ascertain that [the] a client has the full data item.

34. (Original) The method of claim 33 wherein the challenge step comprises requiring that the client attempting to store a data item provide correct answers to inquiries as to the content of portions of the data item, or inquiries that require knowledge of this content.

35. (Original) The method of claim 34 wherein the data item content on which the challenge is based is selected with a degree of randomness.

36. (Previously Presented) The method of claim 1 wherein depositors use the client to store data items in the repository, and at least some depositors are required to provide identification.

37. (Original) The method of claim 36 wherein rules for when a depositor must provide identification are selected in order to discourage unlawful distribution of access to the data item.

38. (Original) The method of claim 37 wherein there is a greater degree of user identification or a higher likelihood that user identification will be required when the data item being stored by the depositor has been indicated to be shareable with other users.

39. (Original) The method of claim 37 wherein for a class of data items the items may only be shared if the depositor has provided adequate identification.

40. (Original) The method of claim 38 or 39 wherein identity information about the depositor is made available to anyone able to access the data item, to discourage unlawful sharing.

41. (Original) The method of claim 40 wherein the identity information is stored in an encrypted form that the depositor and users subsequently accessing the shared data item can both read.

42. (Original) The method of claim 41 wherein the repository is not able to decrypt the identity information about the depositor.

43. (Original) The method of claim 37 wherein the identity of some users has not been well verified, but restrictions are placed on sharing of data items deposited by such poorly verified users.

44. (Original) The method of claim 43 further comprising limiting access to data items deposited by a poorly verified user.

45. (Original) The method of claim 44 wherein the limited access is provided by limiting the aggregate bandwidth provided for such accesses.

46. (Original) The method of claim 44 wherein the limited access is provided by limiting the number of simultaneous accesses to the data items.

47. (Currently Amended) The method of claim 1 wherein the client has a directory structure for the data items, the data items are stored in the repository, and the directory structure is not evident to maintainers of the repository [maintainers].

48. (Previously Presented) The method of claim 1 wherein the client program using the repository is a mirroring program which determines which data items to deposit in the repository, and wherein that determination is based at least in part on the result of a comparison of digital fingerprints establishing that certain data items are not in the repository.

49. (Currently Amended) The method of claim 48 wherein mirroring software is downloaded to the client using a bootstrap process, wherein a small bootstrap program is downloaded and executed, and the bootstrap program manages download and installation of a further portion [the remainder] of the mirroring software.

50. (Original) The method of claim 48 wherein the default for deciding what data items to mirror is to mirror all or substantially all data items.

51. (Original) The method of claim 48 wherein the mirroring comprises making a determination of which data items need to be transmitted to the repository, and wherein that determination is based primarily on a comparison of digital fingerprints for data items at the client and data items in the repository.

52. (Original) The method of claim 10 wherein the access-authorization credential is determined in part by computing a hash involving elements of the pathname for a file on the client computer.

53. (Original) The method of claim 52 wherein the path name hash is made unique to a client by introducing a reproducible but randomly chosen element into it.

54. (Original) The method of claim 12 wherein a data item is represented as a composite of data-items, and the component data-items are separately deposited in the repository.



55. (Original) The method of claim 54 wherein lists of fingerprints for data-items making up a composite data-item are deposited as an index data item, which can be given an object-name and used for obtaining access to any of the component data-items.

56. (Original) The method of claim 55 wherein a proof-of-deposit is returned for each component deposit, and some or all of the proofs are presented when the index data item is given an object-name.

57. (Original) The method of claim 56 wherein, when transmitting a composite data-item, the client uses fingerprints to avoid retransmitting components following loss of communication.

58. (Original) The method of claim 57 wherein the index data-item is encrypted with a key that is only made available to the repository at the moment of access.

59. (Original) The method of claim 55 wherein an email message is broken up into component items in such a manner that the individual attachments are separate component data-items.

60. (Original) The method of claim 15 wherein the physical location at which information about named-objects is stored is based on access identifiers, to introduce reproducible pseudorandomness into the physical locations of the named-object data.

61. (Canceled).

62. (Previously Presented) The method of claim 1 wherein an access identifier is formed to provide proof of ownership of the data item stored in the repository, the access identifier is formed by producing a one-way hash including item-identifying information chosen by the client

program to identify the data item, and the one-way hash cannot be reversed to permit the repository to discover the identity of the client program or user.

63. (Currently Amended) The method of claim 62 wherein the item-identifying information is associated with [the data item on] the client.

64. (Original) The method of claim 63 wherein the item-identifying information is derived at least in part from the path name of the data item on the client.

65. (Original) The method of claim 62 wherein user-identifying information is provided to the repository as part of the access-authorization credential.

66. (Original) The method of claim 65 wherein at least some access-authorization credentials can be transferred between users without the use of the repository.

67. (Original) The method of claim 65 wherein at least one class of users is not permitted to transfer access using access-authorization credentials.

68-153. (Canceled).

154. (Previously Presented) The method of claim 1 further comprising encrypting the data item using a key derived from the content of the data item.

155. (Previously Presented) The method of claim 1, 9, 10, 11, 22, 33, 36, 47, 48 or 62 further comprising encrypting the data item using a key derived from the content of the data item.

156. (Previously Presented) The method of claim 1 wherein the data items are widely circulated non- electronic media such as books or music, and the method further comprises converting the widely circulated non-electronic media to a standardized electronic version; storing the standardized electronic version as a data item in the repository; promoting the availability of the standardized electronic version to users with the right to have access, whereby the likelihood of the data repository storing multiple, slightly-different electronic versions of the non-electronic media is reduced.

157. (Previously Presented) The method of claim 48 wherein the aforesaid steps of the method provide a mirroring capability for a personal computer, and mirroring software with instructions for carrying out the aforesaid steps is preconfigured on the personal computer upon purchase.

158. (Previously Presented) The method of claim 48 wherein the aforesaid steps of the method provide a mirroring capability for a personal computer, and mirroring software for carrying out the method is initially configured to mirror essentially all data on the user's computer.

159. (Previously Presented) The method of claim 48 wherein the aforesaid steps of the method provide a mirroring capability for a wireless network device.

160-174. (Canceled).

175. (Previously Presented) The method of claim 1 in which different physical locations comprise different hard disk drives.

176. (Previously Presented) The method of claim 1 in which different physical locations comprise different data servers.

177. (Previously Presented) The method of claim 1 wherein the physical locations each comprise one or more different processors.

178. (Currently Amended) The method of claim 1 wherein the [data repository] physical locations comprise[s] physical storage nodes linked by a network.

179. (Previously Presented) The method of claim 1 wherein determining from the digital fingerprint whether a data item identical to the second data item is already stored in the data repository comprises transmitting over the network the digital fingerprint of the second data item rather than the second data item itself.

180. (Previously Presented) The method of claim 1 wherein the first and second client programs are independent programs.

181. (Previously Presented) The method of claim 180 wherein the independent programs are running on separate computers.

182. (Previously Presented) The method of claim 1 wherein the first and second client programs are the same program running at different times.

183. (Previously Presented) The method of claim 1 wherein at least the first client program comprises a file server.

184. (Previously Presented) The method of claim 1 wherein files and directories are named objects within the data repository.

185. (Previously Presented) The method of claim 1 wherein a structured item is split up into a plurality of data items with the divisions occurring at content dependent boundaries.

186. (Previously Presented) The method of claim 185 wherein the structured data item comprises an e-mail message and the content dependent boundaries are the divisions between email attachments.

187. (Previously Presented) The method of claim 1 wherein a plurality of clients each of which has initiated a process to deposit an identical data item all share read access to a single repository data-item.

188. (Previously Presented) The method of claim 187 wherein clients which have not initiated a process for depositing the identical data item do not possess a credential that authorizes them to read the identical data item.

189. (Previously Presented) The method of claim 187 wherein a reference count that reflects the number of clients that share read access to the single repository data-item has transitioned to zero and the storage space associated with the shared repository data-item is reclaimed.